

AISPG™ / PUBLIC WHITE PAPER 02

THE AI SECURITY & GOVERNANCE AUTHORITY

GOVERNING THE AGENTIC ENTERPRISE.

A PUBLIC ARCHITECTURE PERSPECTIVE FOR AI AGENTS

Identity, authority, tools, data and defensible
decisions.

FOUNDER-ORIGINATED PUBLIC PERSPECTIVE

ALI AQEL / AISPG™ / OCTOBER 2026

PUBLIC EDITION — CONCEPTUAL RESEARCH

FIELD NOTE / 01

When software can act, governance must follow authority.

An agentic workflow can interpret intent, select tools, retrieve data, invoke services and trigger actions. The resulting risk cannot be understood solely by looking at the language model. It sits along the chain of identities, delegation, data boundaries, tool permissions and consequential outcomes.

01 Permission is not purpose

A credential that can execute an action does not establish that the action is authorized for this task and context.

02 Delegation changes risk

An agent may use another service, tool or sub-agent. Accountability must survive each delegation boundary.

03 Evidence must follow action

Authorization context, tool calls, data flows, exceptions and oversight decisions need a coherent trace.

THE AISPG VIEW

The AISPG™ public agentic-governance view is an architectural way to ask and connect these questions. It is not a claim of guaranteed prevention, universal coverage or verified runtime enforcement.

FIELD NOTE / 02

The agentic risk surface is a chain, not a component.

Governance should ask how a task becomes an action and what is lost, inherited or expanded along the path. These examples are hypothetical.

01 / DELEGATION

An assistant hands a task to a privileged tool or another agent without a clear boundary on the authority transferred.

02 / TOOLING

A connected tool accepts a high-impact command; technical capability is mistaken for approved business authority.

03 / INFORMATION

Retrieved documents or external content influence the agent's next step beyond the intended trust boundary.

04 / OVERSIGHT

Human review is nominally required but occurs after an irreversible action or without sufficient context.

THE REFRAME

Do not ask whether each tool has a dashboard. Ask whether the enterprise can reconstruct effective authority, scoped obligations and the evidence supporting a consequential decision.

FIELD NOTE / 03

Seven connected governance questions for agents.

These public questions illustrate which boundaries an enterprise should examine. They are not a disclosure of the proprietary AISPG control catalog.

01 Identity

Which human, service and agent identities actually executed the task?

02 Delegated Authority

What permissions traveled between agents, services and tools?

03 Tool Boundaries

Which operations were allowed, denied, rate-limited or reviewed?

04 Data Purpose

What data crossed the workflow and was each use permitted?

05 Decision Rights

Which consequential outcomes required human approval or escalation?

06 Evidence & Exceptions

What can be traced, challenged and revalidated after the action?

07 Cloud & Runtime Context

What deployment, provider and infrastructure boundaries affected operation?

FIELD NOTE / 04

Five checkpoints in an accountable agentic action.

The conceptual control narrative below shows where governance expectations may be tested. It is intentionally technology-neutral and does not expose proprietary enforcement logic.

01**INTENT**

Capture the approved business purpose, requester, scope and acceptable outcome.

02**AUTHORITY**

Resolve effective identities, permissions, tool eligibility and delegated bounds.

03**EXECUTION**

Constrain tool behavior, data use and escalation requirements during action.

04**EVIDENCE**

Preserve a trace connecting actions and decisions with source records and limits.

05**REASSESSMENT**

Review deviations, new versions, exceptions and the continuing right to act.

DESIGN PRINCIPLE

Missing evidence does not become positive proof. Failed tests, exceptions and uncertainty belong in the conclusion - not hidden behind a composite color or score.

FIELD NOTE / 05

An agent with purchasing power: illustrative journey.

Hypothetical example: an internal procurement agent is asked to request a vendor quote and prepare a purchase. Nothing in this paper implies the workflow was deployed or tested by AISPG.

1

01 / LIMIT THE PURPOSE

Allow quote discovery and draft preparation, but do not assume the user also authorized a purchase.

2

02 / VERIFY DELEGATION

Identify the requesting human, agent service identity and the permissions of connected purchasing tools.

3

03 / APPLY DECISION GATES

Require human approval before financial commitment; retain scope, rationale and tool-output context.

4

04 / RETAIN THE RECORD

Record what was requested, proposed, approved, attempted or denied, including exceptions and later changes.

ILLUSTRATIVE ONLY

This is a teaching scenario, not a customer claim, control test, attestation, benchmark or record of AISPG production activity.

FIELD NOTE / 06

Six assurance questions for agentic leadership.

Operational accountability depends on asking questions that join AI engineering, identity, cloud, security, legal and business ownership.

01 Can we identify the human and service authority behind every material agent action?

02 What can each tool do, and what should it be allowed to do for this purpose?

03 How are prompt-injection and untrusted-content paths considered in the workflow?

04 Where do irreversible or regulated decisions require explicit human approval?

05 Can we reconstruct the action from source evidence without assuming perfect telemetry?

06 What invalidates the approval when models, tools, permissions or context change?

START WITH ONE REAL JOURNEY

Start with one tool-using agent. Trace identity, delegation and decision gates; rehearse a blocked action and capture the resulting review record.

FIELD NOTE / 07

Context. Sources. Further inquiry.

AISP^g's interpretation is independent. The following authoritative publications are external context, not endorsement, adoption, certification or validation of AISP^g.

REFERENCE / 01

NIST AI Risk Management Framework 1.0

NIST, 2023. Voluntary risk-governance framework with Govern, Map, Measure and Manage functions.

<https://airc.nist.gov/airmf-resources/airmf/>

REFERENCE / 02

NIST Generative AI Profile

NIST AI 600-1, 2024. A cross-sectoral profile relating generative AI risk to the AI RMF.

<https://doi.org/10.6028/NIST.AI.600-1>

REFERENCE / 03

ISO/IEC 42001:2023

International standard specifying requirements for an AI management system.

<https://www.iso.org/standard/42001>

REFERENCE / 04

OWASP GenAI Security Project

Official open guidance, including the September 2026 LLM risks and agent-control resources.

<https://genai.owasp.org/>

PUBLIC EDITION / DISCLOSURE BOUNDARY

© 2026 Ali Aqel / AISP^g™. This original paper is provided for education and discussion. AISP^g is a privately owned, founder-originated framework, not a regulator or accredited certifier. Conceptual diagrams and questions do not represent live control validation, licensed operating procedures, legal advice or standards conformity. Proprietary control logic, undisclosed reasoning mechanisms, private architecture and sensitive platform verification remain outside this edition.

EXPLORE THE PUBLIC AUTHORITY EXPERIENCE

AISP^g.ai / Research & qualified access